ON THE STABILITY OF DEEP NETWORKS

Raja Giryes & Guillermo Sapiro
Duke University
{raja.giryes, guillermo.sapiro}@duke.edu

Alex M. Bronstein
Tel Aviv University
bron@eng.tau.ac.il

arXiv:1412.5896v3 [stat.ML] 3 Jun 2015

ABSTRACT

In this work we study the properties of deep neural networks (DNN) with random weights. We formally prove that these networks perform a distance-preserving embedding of the data. Based on this we then draw conclusions on the size of the training data and the networks' structure. A longer version of this paper with more results and details can be found in (Giryes et al., 2015). In particular, we formally prove in (Giryes et al., 2015) that DNN with random Gaussian weights perform a distance-preserving embedding of the data, with a special treatment for in-class and out-of-class data.

1 INTRODUCTION

Deep neural nets (DNN) have led to a revolution in the areas of machine learning, audio analysis, and computer vision. Many state-of-the-art results have been achieved using these architectures. In this work we study the properties of these architectures with random weights. We prove that DNN preserve the distances in the data along their layers and that this property allows stably recovering the original data from the features calculated by the network. Our results provide insights into the outstanding empirically observed performance of DNN and the size of the training data.

Our motivation for studying networks with random weights is threefold. First, one of the differences between the networks used two decades ago and state-of-the-art training strategies is the usage of random initialization of the weights. Second, a series of works (Pinto et al., 2009; Saxe et al., 2011; Cox & Pinto, 2011) empirically showed successful DNN learning techniques based on randomization. Third, recent works that studied the optimization aspect in the training of deep networks also have done so via randomization (Saxe et al., 2014; Dauphin et al., 2014; Choromanska et al., 2015).

Bruna et al. (2013) show that the pooling stage in DNN causes a shift invariance property. Bruna et al. (2014) interpret this step as the removal of phase from a complex signal and show how the signal may be recovered after a pooling stage using phase retrieval methods. In this short note, and for presentation purposes, we do not consider the previously studied pooling step, assuming the data to be properly aligned. We focus on the roles of the layers of a linear operation followed by an element-wise non-linear activation function.

2 STABLE EMBEDDING OF A SINGLE LAYER

We assume the input data to belong to a manifold K with Gaussian mean width

$$\omega(K) := E[\sup_{\mathbf{x},\mathbf{y}\in K} \langle \mathbf{g}, \mathbf{x}-\mathbf{y}\rangle], \tag{1}$$

where the expectation is taken over $\mathbf{g}$ with normal i.i.d. elements. In Section 3 we will illustrate this concept and exemplify the results with Gaussian mixture models (GMM).

We say that $f : \mathbb{R} \to \mathbb{R}$ is a semi-truncated linear function if it is linear on some (possibly, semi-infinite) interval and constant outside of it, $f(0) = 0$, $0 < f(x) \leq x, \forall x > 0$ and

$0 \geq f(x) \geq x, \forall x < 0$. The popular rectified linear unit (ReLU), $f(x) = \max(0, x)$, is an example of such a function, while the sigmoid functions satisfy this property approximately. The following theorem shows that each standard DNN layer performs a stable embedding of the data in the Gromov-Hausdorff sense.

Theorem 1 *Let $\mathbf{M}$ be the linear operator applied at the i-th layer, f the non-linear activation function, and $K \subset \mathbb{S}^{n-1}$ the manifold of the input data for the i-th layer. If $\sqrt{m}\mathbf{M} \in \mathbb{R}^{n \times m}$ is a random matrix with i.i.d normally distributed entries with $m = O(\omega(K)^2)$ being the output dimension, and f is a semi-truncated linear function, then with high probability*

$$\|\mathbf{x} - \mathbf{y}\|_2 \simeq d(f(\mathbf{M}\mathbf{x}), f(\mathbf{M}\mathbf{y})), \quad \forall \mathbf{x}, \mathbf{y} \in K, \tag{2}$$

where $d(\cdot, \cdot)$ is a variant of the Hamming distance that treats the positive values in the vectors as ones. This result implies that the metric of the input data is preserved.

The proof follows from (Plan & Vershynin, 2014) and Klartag & Mendelson (2005).

Mahendran & Vedaldi (2014) demonstrate that it is possible to recover the input of DNN from their output. The next result provides a theoretical justification for their observation by showing that it is possible to recover the input of each layer from its output:

Theorem 2 *Under the assumptions of Theorem 1 there exists a program $\mathcal{A}$ such that*

$$\|\mathbf{x} - \mathcal{A}(f(\mathbf{M}\mathbf{x}))\|_2 \leq \epsilon, \tag{3}$$

where $\epsilon = O\left(\frac{\omega(K)}{\sqrt{m}}\right)$.

The proof follows from Plan & Vershynin (2014).

3 Stable Embedding of the Entire Network

In order to show that the entire network produces a stable embedding of its input, we need to show that the Gaussian mean width does not grow significantly as the data propagate through the layers of the network. Instead of bounding the variation of the Gaussian mean width throughout the network, we bound the change in the covering number $N(K, \epsilon)$, i.e., the lowest number of ℓ_2-balls of radius ϵ that cover K. Having the bound on the covering number, we use Dudleys inequality (Ledoux & Talagrand, 1991), $\omega(K) \leq C \int_0^\infty \sqrt{\log N(K, \epsilon)} d\epsilon$, to bound the Gaussian mean width variation, where C is a constant.

Theorem 3 *Under the assumptions of Theorem 1,*

$$N(f(\mathbf{M}K), \epsilon) \leq N\left(K, \frac{\epsilon}{1 + \frac{\omega(K)}{\sqrt{m}}}\right). \tag{4}$$

Proof: We now present a sketch of the proof, deferring the full proof that treats also the Gaussian mean width directly to a longer version of the paper. It is not hard to see that since a non-linear activation function shrinks the data, then it can not increase the size of the covering; therefore we focus on the linear part. Following (Klartag & Mendelson, 2005, Theorem 1.4), we have that the distances in $\mathbf{M}K$ are the same as the ones in K up to a $1 + \frac{\sqrt{m}}{\omega(K)}$ factor. This is sufficient to complete the proof. □

We demonstrate the implication of the above theorem for a GMM, i.e., K consisting of L Gaussians of dimension k in the ℓ_2-ball. For this model $N(K, \epsilon) = L\left(1 + \frac{2}{\epsilon}\right)^k$ for $\epsilon < 1$ and 1 otherwise (see Mendelson et al. (2008)). Therefore we have that $\omega(K) \leq C'\sqrt{k + \log L}$ and that at each layer the Gaussian mean width grows at most with an order of $1 + \frac{\sqrt{k} + \log L}{\sqrt{m}}$. Similar results can be shown for other models of union of subspaces and low dimensional manifolds.

4 How Many Measurements Are Needed to Train the Network

An important question in deep learning is what is the amount of labeled training samples needed at training. Using Sudakov minoration (Ledoux & Talagrand, 1991), one may get an upper bound on the size of an ϵ-net in K. We have demonstrated that networks with random Gaussian weights realize a stable embedding; consequently, if a network is trained using the screening technique by selecting the best among many networks generated with random weights as suggested in Pinto et al. (2009); Saxe et al. (2011); Cox & Pinto (2011), then the number of data points needed to be used in order to guarantee that the network represents all the data is $O(\exp(\omega(K)^2/\epsilon^2))$. Since $\omega(K)^2$ is a proxy for the data dimension (see Plan & Vershynin (2014)), we conclude that the number of training points grows exponentially with the intrinsic dimension of the data.

5 Discussion and Conclusion

We have shown that DNN with random Gaussian weights perform a distance-preserving embedding of the data. This result provides a relationship between the complexity of the input data and the size of the required training set. In addition, it draws a connection between the dimension of the features produced by the network,which still keep the metric information of the original manifold, and the complexity of the data.

Though we have focused here on the case of DNN with linear filters with random Gaussian entries, it is possible to extend our analysis to distributions such as sub-Gaussian, and to random convolutional filters using proof techniques from (Haupt et al., 2010; Saligrama, 2012; Rauhut et al., 2012; Ai et al., 2014). This and the extension to learned DNN will be presented in an extended version of this note.

Acknowledgments: This work is supported by NSF, DoD and ERC StG 335491.

References

Ai, A., Lapanowski, A., Plan, Y., and Vershynin, R. One-bit compressed sensing with non-gaussian measurements. *to appear in Linear Algebra and its Applications*, 2014.

Bruna, J., LeCun, Y., and Szlam, A. Learning stable group invariant representations with convolutional networks. In *ICLR Workshop*, Jan. 2013.

Bruna, J., Szlam, A., and LeCun, Y. Signal recovery from l_p pooling representations. In *Int. Conf. on Machine Learning (ICML)*, 2014.

Choromanska, A., Henaff, M. B., Mathieu, M., Arous, G. Ben, and LeCun, Y. The loss surfaces of multilayer networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.

Cox, D. and Pinto, N. Beyond simple features: A large-scale feature search approach to unconstrained face recognition. In *IEEE International Conference on Automatic Face Gesture Recognition and Workshops (FG)*, pp. 8–15, March 2011.

Dauphin, Y., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

Giryes, R., Saprio, G., and Bronstein, A. M. Deep neural networks with random gaussian weights: A universal classification strategy? *ArXiv*, 2015. URL `http://arxiv.org/abs/1504.08291`.

Haupt, J., Bajwa, W.U., Raz, G., and Nowak, R. Toeplitz compressed sensing matrices with applications to sparse channel estimation. *IEEE Trans. Inf. Theory*, 56(11):5862–5875, Nov. 2010.

Klartag, B. and Mendelson, S. Empirical processes and random projections. *Journal of Functional Analysis*, 225(1):229–245, Aug. 2005.

Ledoux, Michel and Talagrand, Michel. *Probability in Banach Spaces*. Springer-Verlag, 1991.

Mahendran, A. and Vedaldi, A. Understanding deep image representations by inverting them. *ArXiv*, 2014. URL `http://arxiv.org/abs/1412.0035`.

Mendelson, S., Pajor, A., and Tomczak-Jaegermann, N. Uniform uncertainty principle for Bernoulli and sub-Gaussian ensembles. *Constructive Approximation*, 28:277–289, 2008.

Pinto, N., Doukhan, D., DiCarlo, J. J., and Cox, D. D. A high-throughput screening approach to discovering good forms of biologically inspired visual representation. *PLoS Comput Biol*, 5(11): e1000579, 11 2009.

Plan, Y. and Vershynin, R. Dimension reduction by random hyperplane tessellations. *Discrete and Computational Geometry*, 51(2):438–461, 2014.

Rauhut, H., Romberg, J., and Tropp, J. A. Restricted isometries for partial random circulant matrices. *Appl. Comput. Harmon. Anal.*, 32(2):242–254, Mar. 2012.

Saligrama, V. Aperiodic sequences with uniformly decaying correlations with applications to compressed sensing and system identification. *IEEE Trans. Inf. Theory*, 58(9):6023–6036, Sept. 2012.

Saxe, A., Koh, P. W., Chen, Z., Bhand, M., Suresh, B., and Ng, A. Y. On random weights and unsupervised feature learning. In *Int. Conf. on Machine Learning (ICML)*, pp. 1089–1096, 2011.

Saxe, A., McClelland, J., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural network. In *International Conference on Learning Representations (ICLR)*, 2014.